Porting NOvA to OSG: Lessons Learned

Gabriele Garzoglio

Oct 23, 2013

NOvA on OSG – Proof-of-principle

- 1,000,000 ev generated with ~10k jobs for 88,535 CPU h in 2 weeks of ops and 2 TB of data
- Run on OSG at SMU (dedicated), UNL, Uwisc, UC, UCSD and O(100) jobs at FermiCloud
- Operations consisted of 1 submission + 2 recoveries; done through SAM
- Spent about 10% more resources than expected due to preemption

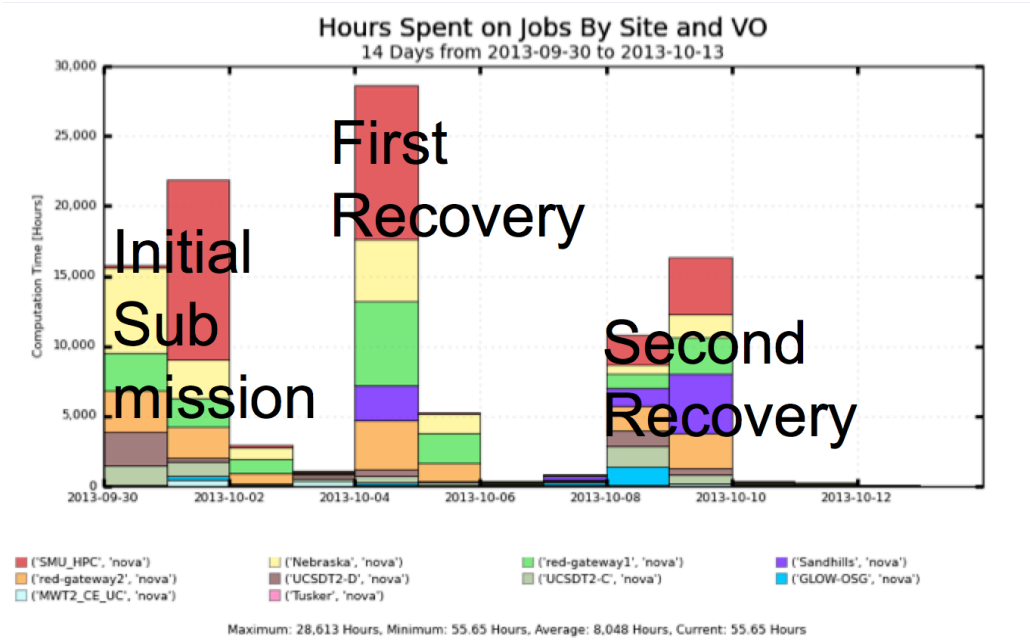

- Now planning to generate 16M ev, running 160k jobs (16x10k) for 800k CPU h and 32 TB of output
- Counting 1,500 CPU DC expect 2 months of ops.
- Adding dedicated resources at OSC & Harvard + Clouds at AWS & Notre Dame

Observations

- Portability of NOvA code started several months ago.
  - Tests of CVMFS were already successful at SMU and Harvard
  - The code was essentially portable in Sep.
- Coordinated push started in Sep.
  - Lot of work done before that on various thrusts

Lessons: Governance

- Follow a defined process for the porting activity
  - The current effort was started without any formal process
  - We lacked coordination, clear communication channels, agreed documentation outlets, appropriate integration environments, well-understood test cases
  - We have started to write up the "Onboarding new communities – best practices"
- Understand and follow the priorities of the community doing the porting
- Be candid about the amount of effort and time required to do the porting
- Be flexible on goals set by the community
  - Benchmarks of the scientific app may not be full understood
    - For NOvA, estimated generation at 30 sec / ev → measured 3.5 – 5 min /ev

Lessons: Communication

- Organize the structure of the communication channels
  - For NOvA, almost all of the application porting was already done
  - We gathered the appropriate SCD and NOvA management and technical experts in a single mailing list
    - Most of the communication was technical
  - We organized daily stand up meetings for 1 month, for the final sprint
    - Do stand up meeting in the afternoon
    - Beware of meeting fatigue
- For end-to-end future efforts…
  - Separate communication in categories (technical, informational, etc.) and associate stakeholders to different mailing lists
  - Meeting frequency may also need to follow the “sprints”
  - Establish IM chat channels for technical communication
  - Establish better policies to associate work to tickets
  - Improve / automate weekly status reports with tickets

Lessons: Integration Environments

- Turn-around time to integrate new features was slow
  - We should provide interactive WN-like environment to test the application porting and new releases of the integrated software (wrappers, IFDH, etc.)
  - This can be an "on-demand" service (a VM available when needed)
- Site on-boarding was *ad-hoc* and slow
  - Define test cases and integration conditions more clearly. Test cases developed this time:
    - "Hello new site", OSG Factory integration, basic tools sanity (SRM, CVMFS), NOvA short test job (to go to production), regular NOvA job

Lessons: Operations

- Set expectations for successful porting activities
  - This takes effort from at least one experimenter and the SCD support team
  - Requires intensive turn around until operational patterns are well established
  - Transfer of knowledge from the experimenter to the experiment operational team requires training

Lessons: Technical

- Provide better integration of Grid and Cloud accounting
- Improve scalability of submission infrastructure → Client / Server JobSub
- Improve resource groups organization…
  - FNAL, OSG dedicated, OSG opportunistic, private clouds, per-pay clouds, …
- … and provisioning policies
  - go to OSG for MC, to FNAL for reco, to per-pay clouds upon request, …
- Improve data transfer mechanisms
  - Information system on storage; integration with job management; direct vs. local output staging; …